arXiv:1404.4911v3 [math.OC] 20 Sep 2015

Communication Delay Co-Design in $\mathcal{H}_2$ Distributed Control Using Atomic Norm Minimization

Nikolai Matni [*]

September 22, 2015

Abstract

When designing distributed controllers for large-scale systems, the actuation, sensing and communication architectures of the controller can no longer be taken as given. In particular, controllers implemented using dense architectures typically outperform controllers implemented using simpler ones – however, it is also desirable to minimize the cost of building the architecture used to implement a controller. The recently introduced Regularization for Design (RFD) framework poses the controller architecture/control law co-design problem as one of jointly optimizing the competing metrics of controller architecture cost and closed loop performance, and shows that this task can be accomplished by augmenting the variational solution to an optimal control problem with a suitable atomic norm penalty. Although explicit constructions for atomic norms useful for the design of actuation, sensing and joint actuation/sensing architectures are introduced, no such construction is given for atomic norms used to design communication architectures. This paper describes an atomic norm that can be used to design communication architectures for which the resulting distributed optimal controller is specified by the solution to a convex program. Using this atomic norm we then show that in the context of $\mathcal{H}_2$ distributed optimal control, the communication architecture/control law co-design task can be performed through the use of finite dimensional second order cone programming.

1 Introduction

Large-scale systems represent an important class of application areas for the control engineer – prominent examples include the smart-grid, software defined networking (SDN) and automated highways. For such large-scale systems, designing the controller *architecture* – placing sensors and actuators as well as the communication links between them – is now also an important part of the controller synthesis process. Indeed controllers with denser actuation, sensing and communication architectures will typically outperform those with simpler architectures – however it is also desirable to minimize the cost of constructing a controller architecture.

In [2], the author of this paper and V. Chandrasekaran address the problem of jointly optimizing the architectural complexity of a distributed optimal controller and the closed loop performance that it achieves by introducing the Regularization for Design (RFD) framework. In RFD, controllers with complicated architectures are viewed as being composed of atomic controllers with simpler

[*]N. Matni is with the Department of Control and Dynamical Systems, California Institute of Technology. 1200 E California Blvd., Pasadena, CA, 91125. (626) 395-6247. nmatni@caltech.edu.

[†]This research was in part supported by the NSF, AFOSR, ARPA-E, and the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. A preliminary version of this work [1] has appeared at the 52nd Annual Conference on Decision and Control in December, 2013.

architectures – this family of simple controllers is then used to construct various *atomic norms* [3–5] that penalize the use of specific architectural resources, such as actuators, sensors or additional communication links. These atomic norms are then added as a penalty function to the variational solution to an optimal control problem (formulated in the model matching framework), allowing the controller designer to explore the tradeoff between architectural complexity and closed loop performance by varying the weight on the atomic norm penalty in the resulting convex optimization problem.

In [2] we give explicit constructions of atomic norms useful for the design of actuation, sensing and joint actuation/sensing architectures, but do not address how to construct an atomic norm for communication architecture design. Indeed constructing a suitable atomic norm for communication architecture design has substantial technical challenges that do not arise in actuation and sensing architecture design: we address these challenges in this paper. We model a distributed controller as a collection of sub-controllers, each equipped with a set of actuators and sensors, that exchange their respective measurements with each other subject to *communication delays* imposed by an underlying communication graph. Keeping with the philosophy adopted in RFD [2], we view dense communication architectures, i.e., ones with a large number of communication links between sub-controllers, as being composed of multiple simple *atomic* communication architectures, i.e., ones with a small number of communication links between sub-controllers. Thus the problem of controller communication architecture/control law co-design can be framed as the joint optimization of a suitably defined measure of the communication complexity of the distributed controller and its closed loop performance, in which these two competing metrics are traded off against each other in a principled manner.

In general one can select communication architectures that range in complexity from completely decentralized, i.e., distributed controllers with no communication allowed between sub-controllers, to essentially centralized and without delay, i.e., distributed controllers with instantaneous communication allowed between all sub-controllers. However, if we ask that the distributed optimal controller restricted to the designed communication architecture be specified by the solution to a convex optimization problem then this limits the simplicity of the designed communication scheme [6–9]. In particular a sufficient, and under mild assumptions necessary, condition for a distributed optimal controller to be specified by the solution to a convex optimization problem[1] is that the communication architecture allow sub-controllers to communicate with each other as quickly as their control actions propagate through the plant [8]. Although this condition may seem restrictive, it can often be met in practice by constructing a communication topology that mimics or is a superset of the physical topology of the plant. For example, these delay based conditions may be satisfied in a smart-grid setting if fiber-optic cables are laid down in parallel to transmission lines; in a SDN setting if control packets are given priority in routing protocols; and in an automated highway system setting if vehicles are allowed to communicate wirelessly with nearby vehicles.

When the aforementioned delay based condition is satisfied by a distributed constraint, it is said to be *quadratically invariant* (QI) [7,8]. While the resulting distributed optimal control problem is convex when quadratic invariance holds, it may still be infinite dimensional. Recently it has been shown that in the case of $\mathcal{H}_2$ distributed optimal control subject to QI constraints imposed by a strongly connected communication architecture, i.e. one in which every sub-controller can exchange information with every other sub-controller subject to delay, the resulting distributed optimal controller synthesis problem can be reduced to a finite dimensional convex program, and hence admits an efficient solution [12, 13].[2] In light of these observations, we look to design *strongly connected*

[1]For a more detailed overview of the relationship between information exchange constraints and the convexity of distributed optimal control problems, we refer the reader to [7, 8, 10, 11] and the references therein.

[2]Other solutions exist to the $\mathcal{H}_2$ distributed control problem subject to delay constraints – we refer the reader to

communication architectures that induce *QI* constraint sets – once such a communication architecture is obtained, the methods from [12, 13] can then be used to compute the optimal distributed controller restricted to that communication architecture exactly.

Related prior work: Regularization techniques based on atomic norms have been employed to great success in system identification [14–17]. As far as we are aware, the first instance of the use of regularization for the purpose of designing the architecture of a controller can be found in [18] (these methods were then further developed in [19]), in which an ℓ_1 penalty is used with non-convex optimization to synthesize sparse static state feedback controllers with respect to an $\mathcal{H}_2$ performance metric. Other representative examples include the use of ℓ_1 regularization to design sparse treatment therapies [20], consensus [21, 22] and synchronization [23] topologies; and the use of group norm like penalties to design actuation/sensing schemes [24–26].

Contributions: We show that the communication complexity of a distributed controller can be inferred from the structure of its impulse response elements. We use this observation to provide an explicit construction of an atomic norm [3–5], which we call the communication link norm, that can be incorporated into the RFD framework [2] to design strongly connected communication graphs that generate QI subspaces. As argued above, these two structural properties allow for the distributed optimal controller implemented using the designed communication architecture to be specified by the solution to a finite dimensional convex optimization problem [12, 13]. We also show that by augmenting the variational solution to the $\mathcal{H}_2$ distributed optimal control problem presented in [12, 13] with the communication link norm as a regularizer, the communication architecture/control law co-design problem can be formulated as a second order cone program. By varying the weight on the communication link norm penalty function, the controller designer can use our co-design algorithm to explore the tradeoff between communication architecture complexity and closed loop performance in a principled way via convex optimization. We use these results to formulate a communication architecture/control law co-design algorithm that yields a distributed optimal controller and the communication architecture on which it is to be implemented.

Paper Organization: In §2 we introduce necessary operator theoretic concepts and establish notation. In §3 we formulate the communication architecture/control law co-design problem as the joint optimization of a suitably defined measure of the communication complexity of a distributed controller and the closed loop performance that it achieves. In §4, we show how communication graphs can be used to generate distributed constraints, and show that if a communication graph that generates a QI subspace is augmented with additional communication links, the subspace generated by the resulting communication graph is also QI. We use this observation and techniques from structured linear inverse problems [3] in §5 to construct a convex regularizer that penalizes the use of additional communication links by a distributed controller, and formulate the co-design procedure. In §6 we discuss the computational complexity of the co-design procedure and illustrate the usefulness of our approach with two numerical examples. We end with a discussion in §7.

2 Preliminaries

2.1 Operator Theoretic Preliminaries

We use standard definitions of the Hardy spaces $\mathcal{H}_2$ and $\mathcal{H}_\infty$. We denote the restrictions of $\mathcal{H}_\infty$ and $\mathcal{H}_2$ to the space of real rational proper transfer matrices $\mathcal{R}_p$ by $\mathcal{RH}_\infty$ and $\mathcal{RH}_2$, respectively. As we work in discrete time, the two spaces are equal, and as a matter of convention we refer to this space as $\mathcal{RH}_\infty$. We refer the reader to [27] for a review of this standard material. For a signal

the discussion and references in [13] for a more extensive overview of this literature.

$\mathbf{f} = (f^{(t)})_{t=0}^{\infty}$, we use $\mathbf{f}^{\leq d}$ to denote the truncation of $\mathbf{f}$ to its elements $f^{(t)}$ satisfying $t \leq d$, i.e., $\mathbf{f}^{\leq d} := (f^{(t)})_{t=0}^{d}$. We extend the Banach space ℓ_2^n to the space

$$\ell_{2,e}^n := \{\mathbf{f} : \mathbb{Z}_+ \to \mathbb{R}^n \,|\, \mathbf{f}^{\leq d} \in \ell_2^n \text{ for all } d \in \mathbb{Z}_+\}, \tag{1}$$

where $\mathbb{Z}_+$ ($\mathbb{Z}_{++}$) denotes the set of non-negative (positive) integers. A plant $G \in \mathcal{R}_p^{m\times n}$ can then be viewed as a linear map from $\ell_{2,e}^n$ to $\ell_{2,e}^m$. Unless required, we do not explicitly denote dimensions and we assume that all vectors, operators and spaces are of compatible dimension throughout.

2.2 Notation

We denote elements of $\ell_{2,e}$ with boldface lower case Latin letters, elements of $\mathcal{R}_p$ (which include matrices) with upper case Latin letters, and affine maps from $\mathcal{RH}_\infty$ to $\mathcal{RH}_\infty$ with upper case Fraktur letters such as $\mathfrak{M}$. We denote temporal indices, horizons and delays by lower case Latin letters.

We denote the elements of the power series expansion of a map $G \in \mathcal{RH}_\infty$ by $G^{(t)}$, i.e., $G = \sum_{t=0}^{\infty} \frac{1}{z^t} G^{(t)}$. We use $\mathcal{RH}_\infty^{\leq d}$ to denote the subspace of $\mathcal{RH}_\infty$ composed of finite impulse response (FIR) transfer matrices of horizon d, i.e., $\mathcal{RH}_\infty^{\leq d} := \{G \in \mathcal{RH}_\infty \,|\, G = \sum_{t=0}^{d} \frac{1}{z^t} G^{(t)}\}$. Similarly, we use $\mathcal{RH}_\infty^{\geq d+1}$ to denote the subspace of $\mathcal{RH}_\infty$ composed of transfer matrices with power series expansion elements satisfying $G^{(t)} = 0$ for all $t \leq d$, i.e., $\mathcal{RH}_\infty^{\geq d+1} := \{G \in \mathcal{RH}_\infty \,|\, G = \sum_{t=d+1}^{\infty} \frac{1}{z^t} G^{(t)}\}$. For an element $G \in \mathcal{RH}_\infty$, we use $G^{\leq d}$ to denote the projection of G onto $\mathcal{RH}_\infty^{\leq d}$, and $G^{\geq d+1}$ to denote the projection of G onto $\mathcal{RH}_\infty^{\geq d+1}$, i.e., $G^{\leq d} = \sum_{t=0}^{d} \frac{1}{z^t} G^{(t)}$ and $G^{\geq d+1} = \sum_{t=d+1}^{\infty} \frac{1}{z^t} G^{(t)}$.

Sets are denoted by upper case script letters, such as $\mathscr{S}$, whereas subspaces of an inner product space are denoted by upper case calligraphic letters, such as $\mathcal{S}$. We denote the orthogonal complement of $\mathcal{S}$ with respect to the standard inner product on $\mathcal{RH}_2$ by $\mathcal{S}^\perp$. We use the greek letter Γ to denote the adjacency matrix of a graph, and use labels in the subscript to distinguish among different graphs, i.e., Γ_{base} and Γ_1 correspond to different graphs labeled "base" and "1." We use E_{ij} to denote the matrix with (i,j)th element set to 1 and all others set to 0. We use I_n and 0_n to denote the $n \times n$ dimensional identity matrix and all zeros matrix, respectively. For a p by q block row by block column transfer matrix M partitioned as $M = (M_{ij})$, we define the block support $\text{bsupp}\,(M)$ of the transfer matrix M to be the p by q integer matrix with (i,j)th element set to 1 if M_{ij} is nonzero, and 0 otherwise. Finally, we use the $\star$ superscript to denote that a parameter is the solution to an optimization problem.

3 Communication Architecture Co-Design

In this section we formulate the communication architecture/control law co-design problem as the joint optimization of a suitably defined measure of the communication complexity of the distributed controller and its closed loop performance. In particular, we introduce the convex optimization based solution to the $\mathcal{H}_2$ distributed optimal control problem subject to delays presented in [12,13], and modify this method to perform the communication architecture/control law co-design task.

3.1 Distributed $\mathcal{H}_2$ Optimal Control subject to Delays

To review the relevant results of [12,13], we introduce the discrete-time generalized plant G given by

$$G = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & 0 & D_{12} \\ C_2 & D_{21} & 0 \end{array}\right] = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} \tag{2}$$

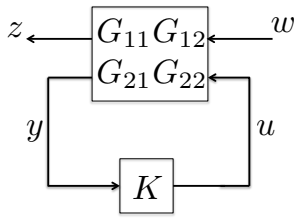

Figure 1: A diagram of the generalized plant defined in (2).

with inputs of dimension p_1, p_2 and outputs of dimension q_1, q_2. As illustrated in Figure 1, this system describes the four transfer matrices from the disturbance and control inputs $\mathbf{w}$ and $\mathbf{u}$, respectively, to the controlled and measured outputs $\mathbf{z}$ and $\mathbf{y}$, respectively. In order to ensure the existence of solutions to the necessary Riccati equations and to obtain simpler formulas, we assume that (A, B_1, C_1) and (A, B_2, C_2) are both stabilizable and detectable, and that

$$D_{12}^\top D_{12} = I,\ D_{21} D_{21}^\top = I,\ C_1^\top D_{12} = 0,\ B_1 D_{21}^\top = 0. \tag{3}$$

Let $\mathcal{S}$ be a subspace that encodes the distributed constraints imposed on the controller K. For example, when some sub-controllers cannot access the measurements of other sub-controllers, the subspace $\mathcal{S}$ enforces corresponding sparsity constraints on the controller K. Alternatively, when sub-controllers can only gain access to other sub-controllers' measurements after a given delay, the subspace $\mathcal{S}$ enforces corresponding delay constraints on the controller K.

The distributed $\mathcal{H}_2$ optimal control problem with subspace constraint $\mathcal{S}$ is then given by

$$\begin{array}{ll} \underset{K \in \mathcal{R}_p}{\text{minimize}} & \left\| G_{11} - G_{12} K (I - G_{22} K)^{-1} G_{21} \right\|_{\mathcal{H}_2}^2 \\ \text{s.t.} & K \in \mathcal{S} \\ & K \text{ internally stabilizes } G \end{array} \tag{4}$$

where the objective function measures the $\mathcal{H}_2$ norm of the closed loop transfer function from the exogenous disturbance $\mathbf{w}$ to the controlled output $\mathbf{z}$, and the first constraint ensures that the controller K respects the distributed constraints imposed by the subspace $\mathcal{S}$.

Optimization problem (4) is in general both infinite dimensional and non-convex. In [12,13], the authors provide an exact and computationally tractable solution to optimization problem (4) when the distributed constraint $\mathcal{S}$ is *QI* [7] with respect to G_{22}[3] and is generated by a *strongly connected* communication graph. We say that a distributed constraint $\mathcal{S}$ is generated by a strongly connected communication graph[4] if it admits a decomposition of the form

$$\mathcal{S} = \mathcal{F} \oplus \frac{1}{z^{d+1}} \mathcal{R}_p,\ \mathcal{F} = \oplus_{t=1}^{d} \frac{1}{z^t} \mathcal{F}^{(t)} \tag{5}$$

for some positive integer d, and some subspaces $\mathcal{F}^{(t)} \subset \mathbb{R}^{p_2 \times q_2}$. In §4 we show how a strongly connected communication graph between sub-controllers can be used to define a subspace $\mathcal{S}$ that admits a decomposition (5).

[3] A subspace $\mathcal{S}$ is said to be QI with respect to G_{22} if $KG_{22}K \in \mathcal{S}$ for all $K \in \mathcal{S}$. When quadratic invariance holds, we have that $K \in \mathcal{S}$ if and only if $K(I - G_{22}K)^{-1} \in \mathcal{S}$; this key property allows for the convex parameterization (6) of the distributed optimal control problem (4).

[4] We consider subspaces $\mathcal{S}$ that are strictly proper so that the reader can use the exact results presented in [13]. The authors of [13] do however note that their method extends to non-strictly proper controllers at the expense of more complicated formulas.

Restricting ourselves to distributed constraints $\mathcal{S}$ that are QI with respect to G_{22} and that admit a decomposition of the form (5) allows us to pose the optimal control problem (4) as the following convex model matching problem

$$\begin{array}{ll} \underset{Q \in \mathcal{RH}_\infty}{\text{minimize}} & \|P_{11} - P_{12} Q P_{21}\|_{\mathcal{H}_2}^2 \\ \text{s.t.} & \mathfrak{C}\left(Q^{\leq d}\right) \in \mathcal{F} \end{array} \tag{6}$$

through the use of a suitable Youla parameterization, where the $P_{ij} \in \mathcal{RH}_\infty$ are appropriately defined stable transfer matrices and $\mathfrak{C} : \mathcal{RH}_\infty^{\leq d} \to \mathcal{RH}_\infty^{\leq d}$ is an appropriately defined affine map (cf. §III-B of [13]). It is further shown in [13] that the solution $Q^\star$ to the distributed model matching problem (6) with QI constraint $\mathcal{S}$ admitting decomposition (5) is specified in terms of the solution to a finite dimensional convex quadratic program.

Theorem 1 (Theorem 3 in [13]) *Let $\mathcal{S}$ be QI under G_{22} and admit a decomposition as in (5). Let $Q^\star \in \mathcal{S} \cap \mathcal{RH}_\infty$ be the optimal solution to the convex model matching problem (6). Then $(Q^\star)^{\geq d+1} = 0$ and*

$$(Q^\star)^{\leq d} = \underset{V \in \mathcal{RH}_\infty^{\leq d}}{\arg\min} \quad \|\mathfrak{L}(V)\|_{\mathcal{H}_2}^2 \;\; s.t. \;\; \mathfrak{C}(V) \in \mathcal{F}, \tag{7}$$

where $\mathfrak{L}$ is a linear map from $\mathcal{RH}_\infty^{\leq d}$ to $\mathcal{RH}_\infty^{\leq d}$, and $\mathfrak{C}$ is the affine map from $\mathcal{RH}_\infty^{\leq d}$ to $\mathcal{RH}_\infty^{\leq d}$ used to specify the model matching problem (6). Furthermore, the optimal cost achieved by $Q^\star$ in the optimization problem (6) is given by

$$\|P_{11}\|_{\mathcal{H}_2}^2 + \left\|\mathfrak{L}\left((Q^\star)^{\leq d}\right)\right\|_{\mathcal{H}_2}^2. \tag{8}$$

Remark 1 *The term $\left\|\mathfrak{L}\left((Q^\star)^{\leq d}\right)\right\|_{\mathcal{H}_2}^2$ in the optimal cost (8) quantifies the deviation of the performance achieved by the distributed optimal controller from that achieved by the centralized optimal controller.*

The optimization problem (7) is finite dimensional because the maps $\mathfrak{L}$ and $\mathfrak{C}$ are both finite dimensional (they map the finite dimensional space $\mathcal{RH}_\infty^{\leq d}$ into itself) and act on the finite dimensional transfer matrix $V \in \mathcal{RH}_\infty^{\leq d}$. These maps can be computed in terms of the state-space parameters of the generalized plant (2) and the solution to appropriate Riccati equations (cf. §III-B and §IV-A of [13]). Under the assumptions (3) the map $\mathfrak{L}$ is injective, and hence the convex quadratic program (7) has a unique optimal solution $(Q^\star)^{\leq d}$.

As the distributed constraint $\mathcal{S}$ is assumed to be QI, the optimal distributed controller $K^\star \in \mathcal{S}$ specified by the solution to the non-convex optimization problem (4) can be recovered from the optimal Youla parameter $Q^\star \in \mathcal{S}$ through a suitable linear fractional transformation (cf. Theorem 3 of [13]).

Remark 2 *If the state-space matrix A specified in the generalized plant (2) is of dimension $s \times s$, then the resulting optimal controller $K^\star$ admits a state-space realization of order $s + q_2 d$. As argued in [13], this is at worst within a constant factor of the minimal realization order.*

3.2 Communication Delay Co-Design via Convex Optimization

Although our objective is to design the communication graph on which the distributed controller K is implemented, for the computational reasons described in §3.1 it is preferable to solve a problem in terms of the Youla parameter Q as this leads to the convex optimization problems (6) and (7). In

order to perform the communication architecture/control law co-design task in the Youla domain, we restrict ourselves to designing strongly connected communication architectures that generate QI subspaces, i.e., subspaces that are QI and that admit a decomposition of the form (5). As argued in §1, this is a practically relevant class of communication architectures to consider, and further, based on the previous discussion it is then possible to solve for the resulting distributed optimal controller restricted to the designed communication architecture using the results of Theorem 1.

Our approach to accomplish the co-design task is to remove the subspace constraint $\mathfrak{C}(V) \in \mathcal{F}$, which encodes the distributed structure of the controller, from the optimization problem (7) and to augment the objective of the optimization problem with a convex penalty function that instead induces suitable structure in $\mathfrak{C}(V)$. In particular, we seek a convex penalty function $\|\cdot\|_{\text{comm}}$ and horizon d such that the structure of $\mathfrak{C}(V^\star)$, where $V^\star$ is the solution to

$$\underset{V \in \mathcal{RH}_\infty^{\leq d}}{\text{minimize}} \quad \|\mathfrak{L}(V)\|_{\mathcal{H}_2}^2 + \lambda \|\mathfrak{C}(V)\|_{\text{comm}}, \tag{9}$$

can be used to define an appropriate QI subspace $\mathcal{S}$ that admits a decomposition of the form (5). Imposing that the designed subspace $\mathcal{S}$ be QI ensures that the structure induced in $\mathfrak{C}(V^\star)$ corresponds to the structure of the resulting distributed controller $K^\star$. Further imposing that the designed subspace $\mathcal{S}$ admit a decomposition of the form (5) ensures that the distributed optimal controller restricted to lie in the subspace $\mathcal{S}$ can be computed using Theorem 1.

Remark 3 *The regularization weight $\lambda \geq 0$ allows the controller designer to tradeoff between closed loop performance (as measured by $\|\mathfrak{L}(V)\|_{\mathcal{H}_2}^2$) and communication complexity (as measured by $\|\mathfrak{C}(V)\|_{comm}$).*

In order to define an appropriate convex penalty $\|\cdot\|_{\text{comm}}$, we need to understand how a communication graph between sub-controllers defines the subspace $\mathcal{F}$ in which $\mathfrak{C}(V)$ is constrained to lie in optimization problem (7) – this in turn informs what structure to induce in $\mathfrak{C}(V^\star)$ in the regularized optimization problem (9). To that end, in §4 we define a simple communication protocol between sub-controllers that allows communication graphs to be associated with distributed subspace constraints in a natural way. Within this framework, we show that if a communication graph generates a distributed subspace $\mathcal{S}$ that is QI with respect to G_{22}, then adding additional communication links to this graph preserves the QI property of the distributed subspace that it generates. We use this observation to pose the communication architecture design problem as one of augmenting a suitably defined base communication graph, namely a simple graph that generates a QI subspace, with additional communication links.

4 Communication Graphs and Quadratically Invariant Subspaces

This section first shows how a communication graph connecting sub-controllers can be used to define the subspace $\mathcal{S}$ in which the controller K is constrained to lie in the distributed optimal control problem (4). In particular, if two sub-controllers exchange information using the shortest path between them on an underlying communication graph, then there is a natural way of generating a subspace constraint from the adjacency matrix of that graph. Under this information exchange protocol, we then define a set of strongly connected communication graphs that generate subspace constraints that are QI with respect to a plant G_{22} in terms of a *base* and a *maximal* communication graph. This approach allows the controller designer to specify which communication links between sub-controllers are *physically realizable*, i.e., which communication links can be built subject to the physical constraints of the system.

4.1 Generating Subspaces from Communication Graphs

Consider a generalized plant (2) comprised of n sub-plants, each equipped with its own sub-controller. Let $\mathscr{N} := \{1, \dots, n\}$ and label each sub-controller by a number $i \in \mathscr{N}$. To each such sub-controller i associate a space of possible control actions $\mathcal{U}_i = \ell_{2,e}^{p_{2,i}}$ and a space of possible output measurements $\mathcal{Y}_i = \ell_{2,e}^{q_{2,i}}$, and define the overall control and measurement spaces as $\mathcal{U} := \mathcal{U}_1 \times \cdots \times \mathcal{U}_n$ and $\mathcal{Y} := \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_n$, respectively.

Then, for any pair of sub-controllers i and j, the $(i,j)^{\text{th}}$ block of G_{22} is the mapping from the control action $\mathbf{u}_j$ taken by sub-controller j to the measurement $\mathbf{y}_i$ of sub-controller i, i.e., $(G_{22})_{ij} : \mathcal{U}_j \to \mathcal{Y}_i$. Similarly, the mapping from the measurement $\mathbf{y}_j$, transmitted by sub-controller j, to the control action $\mathbf{u}_i$ taken by sub-controller i is given by $K_{ij} : \mathcal{Y}_j \to \mathcal{U}_i$.

We then form the overall measurement and control vectors

$$\mathbf{y} = \begin{bmatrix} (\mathbf{y}_1)^\top & \cdots & (\mathbf{y}_n)^\top \end{bmatrix}^\top, \qquad \mathbf{u} = \begin{bmatrix} (\mathbf{u}_1)^\top & \cdots & (\mathbf{u}_n)^\top \end{bmatrix}^\top \tag{10}$$

leading to the natural block-wise partitions of the plant G_{22}

$$G_{22} = \begin{bmatrix} (G_{22})_{11} & \cdots & (G_{22})_{1n} \\ \vdots & \ddots & \vdots \\ (G_{22})_{n1} & \cdots & (G_{22})_{nn} \end{bmatrix} \tag{11}$$

and of the controller K

$$K = \begin{bmatrix} K_{11} & \cdots & K_{1n} \\ \vdots & \ddots & \vdots \\ K_{n1} & \cdots & K_{nn} \end{bmatrix}. \tag{12}$$

We assume that sub-controllers exchange measurements with each other subject to delays imposed by an underlying communication graph – specifically, we assume that sub-controller i has access to sub-controller j's measurement $\mathbf{y}_j$ with delay specified by the length of the shortest path from sub-controller j to sub-controller i in the communication graph. Formally, let Γ be the adjacency matrix of the communication graph between sub-controllers, i.e., Γ is the integer matrix with rows and columns indexed by $\mathscr{N}$, such that Γ_{kl} is equal to 1 if there is an edge from l to k, and 0 otherwise. The *communication delay* from sub-controller j to sub-controller i is then given by the length of the shortest path from j to i as specified by the adjacency matrix gamma Γ. In particular, we define[5] the communication delay from sub-controller j to sub-controller i to be given by

$$c_{ij} := \min \left\{ d \in \mathbb{Z}_+ \,\middle|\, \Gamma_{ij}^d \neq 0 \right\} \tag{13}$$

if an integer satisfying the condition in (13) exists, and set $c_{ij} = \infty$ otherwise.

We say that a strictly proper distributed controller K can be *implemented* on a communication graph with adjacency matrix Γ if for all $i, j \in \mathscr{N}$, we have that the the (i,j)th block of the controller K satisfies $K_{ij}^{(t)} = 0$ for all positive integers $t \leq c_{ij}$, or equivalently, that $K_{ij} \in \frac{1}{z^{c_{ij}+1}} \mathcal{R}_p$. In words, this says that sub-controller j only has access to the measurement $\mathbf{y}_i$ from sub-controller i after c_{ij} time steps, the length of the shortest path from j to i in the communication graph, and can only take actions based on this measurement after a computational delay of one time step.[6] More succinctly, this condition holds if $\text{bsupp}\left(K^{(t)}\right) \subseteq \text{bsupp}\left(\Gamma^{t-1}\right)$ for all $t \geq 1$.

[5] See Lemma 8.1.2 of [28] for a graph theoretic justification of this definition.

[6] This computational delay is included to ensure that the resulting controller is strictly proper.

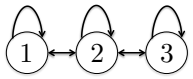

Figure 2: Three subsystem chain example

If Γ is the adjacency matrix of a strongly connected graph, then there exists a path between all ordered pairs of sub-controllers $(i,j) \in \mathscr{N} \times \mathscr{N}$ – this implies that there exists a positive delay $d(\Gamma)$ after which a given measurement $\mathbf{y}_j$ is available to all sub-controllers. In particular, we define the delay $d(\Gamma)$ associated with the adjacency matrix Γ to be

$$d\left(\Gamma\right) := \sup\left\{\tau \in \mathbb{Z}_{++} \,\middle|\, \exists (k,l) \in \mathscr{N} \times \mathscr{N} \text{ s.t. } \Gamma_{kl}^{\tau-1} = 0\right\}. \tag{14}$$

Using this convention all measurements $y_j^{(t)}$ are available to all sub-controllers by time $t + d(\Gamma) + 1$. When the delay $d(\Gamma)$ is finite, we say that Γ is a *strongly connected adjacency matrix*, as it defines a strongly connected communication graph.

We define the subspace $\mathcal{S}(\Gamma)$ generated by a strongly connected adjacency matrix Γ to be

$$\mathcal{S}(\Gamma) := \mathcal{F}(\Gamma) \oplus \frac{1}{z^{d(\Gamma)+1}}\mathcal{R}_p, \tag{15}$$

where $d(\Gamma)$ is as defined in (14), and $\mathcal{F}(\Gamma) := \oplus_{t=1}^{d} \frac{1}{z^t}\mathcal{F}^{(t)}(\Gamma)$ is specified by the subspaces

$$\mathcal{F}^{(t)}(\Gamma) := \left\{M \in \mathbb{R}^{p_2 \times q_2} \,\middle|\, \operatorname{bsupp}\left(M\right) \subseteq \operatorname{bsupp}\left(\Gamma^{t-1}\right)\right\}. \tag{16}$$

It is then immediate that a controller K can be implemented on the communication graph Γ if and only if $K \in \mathcal{S}(\Gamma)$.

Example 1 *Consider the communication graph illustrated in Figure 2 with strongly connected adjacency matrix* $\Gamma_{3\text{-}chain}$ *given by*

$$\Gamma_{3\text{-}chain} = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}. \tag{17}$$

This communication graph generates the subspace

$$\mathcal{S}(\Gamma_{3\text{-}chain}) := \frac{1}{z}\begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix} \oplus \frac{1}{z^2}\begin{bmatrix} * & * & 0 \\ * & * & * \\ 0 & * & * \end{bmatrix} \oplus \frac{1}{z^3}\mathcal{R}_p, \tag{18}$$

where $*$ *is used to denote a space of appropriately sized real matrices. The communication delays associated with this graph are then given by* $c_{ij} = |i-j|$ *(e.g.,* $c_{11} = 0$*,* $c_{12} = 1$ *and* $c_{13} = 2$*). We also have that* $d(\Gamma_{3\text{-}chain}) = 2$*, which is the length of the longest path between nodes in this graph, and that*

$$\mathcal{F}(\Gamma_{3\text{-}chain}) = \frac{1}{z}\begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix} \oplus \frac{1}{z^2}\begin{bmatrix} * & * & 0 \\ * & * & * \\ 0 & * & * \end{bmatrix} \subset \mathcal{RH}_\infty^{\leq 2}.$$

■

Thus, given such a strongly connected adjacency matrix Γ, the distributed optimal controller $K^\star$ implemented using the graph specified by Γ can be obtained by solving the optimization problem (4) with subspace constraint $\mathcal{S}(\Gamma)$ – however, this optimization problem can only be reformulated as the convex programs (6) and (7) if the subspace $\mathcal{S}(\Gamma)$ is QI with respect to G_{22} [9].

4.2 Quadratically Invariant Communication Graphs

The discussion of §3 and §4.1 shows that communication graphs that are strongly connected and that generate a subspace (15) that is QI with respect to G_{22} allow for the distributed optimal control problem (4) to be solved via the finite dimensional convex program (7). In this subsection, we characterize a set of such communication graphs in terms of a *base QI* and a *maximal QI communication graph* corresponding to a plant G_{22}. The base QI communication graph defines a simple communication architecture that generates a QI subspace, whereas the maximal QI communication graph is the densest communication architecture that can be built given the physical constraints of the system.

We assume that the sub-controllers have disjoint measurement and actuation channels, i.e., that B_2 and C_2 are block-diagonal, and that the dynamics of the system are strongly connected, i.e., that $\text{bsupp}\,(A)$ corresponds to the adjacency matrix of a strongly connected graph. We discuss alternative approaches for when these assumptions do not hold in §7. For the sake of brevity, we often refer to a communication graph by its adjacency matrix Γ.

The base QI communication graph

Our objective is to identify a simple communication graph, i.e., a graph defined by a sparse adjacency matrix Γ_{base}, such that the resulting subspace $\mathcal{S}(\Gamma_{\text{base}})$ is QI with respect to G_{22}. To that end, let the *base QI communication graph* of plant G_{22} with realization (2) be specified by the adjacency matrix

$$\Gamma_{\text{base}} := \text{bsupp}\,(A)\,. \tag{19}$$

Notice that under the block-diagonal assumptions imposed on the state-space parameters B_2 and C_2, this implies that Γ_{base} mimics or is a superset of the physical topology of the plant G_{22}, as $\text{bsupp}\left(G_{22}^{(t)}\right) = \text{bsupp}\left(C_2 A^{t-1} B_2\right) \subseteq \text{bsupp}\,(A)^{t-1}$.

Define the *propagation delay* from sub-plant j to sub-plant i of a plant G_{22} to be the largest integer p_{ij} such that

$$(G_{22})_{ij} \in \frac{1}{z^{p_{ij}}}\mathcal{R}_p. \tag{20}$$

It is shown in [8] that if a subspace $\mathcal{S}$ constrains the blocks of the controller K to satisfy $K_{kl} \in \frac{1}{z^{c_{kl}+1}}\mathcal{R}_p$, and the communication delays[7] $\{c_{kl}\}$ satisfy the triangle inequality $c_{ki} + c_{ij} \geq c_{kj}$, then $\mathcal{S}$ is QI with respect to G_{22} if

$$c_{ij} \leq p_{ij} + 1 \tag{21}$$

for all $i, j \in \mathscr{N}$. An intuitive interpretation of this condition is that $\mathcal{S}$ is QI if it allows sub-controllers to communicate with each other as fast as their control actions propagate through the plant. Since we take the base QI communication graph Γ_{base} to mimic the topology of the plant G_{22}, we expect this condition to hold and for $\mathcal{S}(\Gamma_{\text{base}})$ to be QI with respect to G_{22}. We formalize this intuition in the following lemma.

Lemma 1 *Let the plant G_{22} be specified by state-space parameters (A, B_2, C_2), and suppose that B_2 and C_2 are block diagonal. Let $\{p_{ij}\}$ denote the propagation delays of the plant G_{22} as defined in* (20)*. Assume that Γ_{base}, as specified as in equation* (19)*, is a strongly connected adjacency matrix, and let $\{b_{ij}\}$ denote the communication delays* (13) *imposed by the adjacency matrix Γ_{base}. The communication delays $\{b_{ij}\}$ then satisfy condition* (21) *and the subspace $\mathcal{S}(\Gamma_{base})$ is quadratically invariant with respect to G_{22}.*

[7] These are equivalent to the prior definition (13) of communication delays $\{c_{kl}\}$.

Proof: The definition of the base QI communication graph Γ_{base} and the assumption that B_2 and C_2 are block-diagonal imply that $\text{bsupp}\left(G_{22}^{(t)}\right) \subseteq \text{bsupp}\left(A^{t-1}\right) \subseteq \text{bsupp}\left(\Gamma_{\text{base}}^{t-1}\right)$. This in turn can be verified to guarantee that (21) holds. Thus it suffices to show that the communication delays $\{b_{kl}\}$ satisfy the triangle inequality $b_{ki} + b_{ij} \geq b_{kj}$ for all $i, j, k \in \mathscr{N}$. First observe that (i) $b_{ii} + b_{ii} \geq b_{ii}$, and (ii) $b_{ii} + b_{ij} \geq b_{ij}$, as all $b_{ij} \geq 0$. Thus it remains to show that $b_{ki} + b_{ij} \geq b_{kj}$ for $i \neq j \neq k$. Suppose, seeking contradiction, that

$$b_{ki} + b_{ij} < b_{kj}. \tag{22}$$

Note that by definition (13) of the communication delays and Lemma 8.1.2 of [28], the inequality (22) is equivalent to

$$\begin{aligned} &\min\{r \,|\, \exists \text{ path of length } r \text{ from } i \text{ to } k\}+ \\ &\quad \min\{r \,|\, \exists \text{ path of length } r \text{ from } j \text{ to } i\} < \\ &\qquad \min\{r \,|\, \exists \text{ path of length } r \text{ from } j \text{ to } k\}. \end{aligned} \tag{23}$$

Notice however that we must have that

$$\begin{aligned} &\min\{r \,|\, \exists \text{ path of length } r \text{ from } j \text{ to } k\} \leq \\ &\quad \min\{r \,|\, \exists \text{ path of length } r \text{ from } j \text{ to } i\}+ \\ &\qquad \min\{r \,|\, \exists \text{ path of length } r \text{ from } i \text{ to } k\}, \end{aligned} \tag{24}$$

as the concatenation of a path from j to i and a path from i to k yields a path from j to k. Combining inequalities (22) and (24) yields the desired contradiction, proving the result. ∎

Lemma 1 thus provides a simple means of constructing a base QI communication graph by taking a communication topology that mimics the physical topology of the plant G_{22}.

Augmenting the base QI communication graph

The delay condition (21) suggests that a natural way of constructing QI communication architectures given a base QI communication graph is to augment the base graph with additional communication links, as adding a link to a communication graph can only decrease its communication delays c_{ij}.

Proposition 1 *Let Γ_{base} be defined as in* (19)*, and let Γ be an adjacency matrix satisfying $\text{bsupp}\,(\Gamma_{base}) \subset \text{bsupp}\,(\Gamma)$. Then the generated subspace $\mathcal{S}(\Gamma)$, as defined in* (15)*, is quadratically invariant with respect to G_{22}.*

Proof: Let $\{b_{ij}\}$ and $\{c_{ij}\}$ denote the communication delays associated with the base QI communication graph Γ_{base} and the augmented communication graph Γ, respectively. It follows from the definition of the communication delays (13) that the support nesting condition $\text{bsupp}\,(\Gamma_{\text{base}}) \subset \text{bsupp}\,(\Gamma)$ implies that $b_{ij} \geq c_{ij}$ for all $i, j \in \mathscr{N}$. By Lemma 1 we have that $b_{ij} \leq p_{ij} + 1$, and therefore $c_{ij} \leq b_{ij} \leq p_{ij} + 1$. An identical argument to that used to prove Lemma 1 shows that the delays c_{ij} satisfy the required triangle inequality, implying that $\mathcal{S}(\Gamma)$ is QI with respect to G_{22}. ∎

In words, the nesting condition $\text{bsupp}\,(\Gamma_{\text{base}}) \subset \text{bsupp}\,(\Gamma)$ simply means that the communication graph Γ can be constructed by adding communication links to the base QI communication graph Γ_{base}. It follows that any graph built by augmenting Γ_{base} with additional communication links generates a QI subspace (15).

Remark 4 *Although we have suggested a specific construction for Γ_{base}, Proposition 1 makes clear that any strongly connected graph that generates a subspace constraint that is QI with respect to G_{22} can be used as the base QI communication graph. We discuss the implications of this added flexibility in §7.*

The maximal QI communication graph

In order to augment the base QI communication graph in a physically relevant way, one must first specify what additional communication links can be built given the physical constraints of the system. For example, if two sub-controllers are separated by a large physical distance, it may not be possible to build a direct communication link between them. The set of additional communication links that can be physically constructed is application dependent – we therefore assume that the controller designer has specified a collection $\mathscr{E}$ of directed edges that define what communication links can be built in addition to those already present in the base QI communication graph. In particular, we assume that it is possible to build a direct communication link from sub-controller j to sub-controller i, i.e., to build a communication graph $\Gamma_{\text{built}} = \Gamma_{\text{base}} + \Gamma$ with $\Gamma_{ij} = 1$, only if $(i, j) \in \mathscr{E}$.

Given a collection of directed edges $\mathscr{E}$, the *maximal QI communication graph* $\Gamma_{\max}$ is given by

$$\Gamma_{\max} := \Gamma_{\text{base}} + M, \tag{25}$$

where M is a $n \times n$ dimensional matrix with M_{ij} set to 1 if $(i, j) \in \mathscr{E}$ and 0 otherwise. In words, the maximal QI adjacency matrix $\Gamma_{\max}$ specifies a communication graph that uses all possible communication links listed in the set $\mathscr{E}$, in addition to those links already used by the base QI communication graph. Consequently, we say that a communication graph can be *physically built* if its adjacency matrix Γ satisfies

$$\operatorname{bsupp}(\Gamma) \subseteq \operatorname{bsupp}(\Gamma_{\max}), \tag{26}$$

i.e., if it can be built from communication links used by the base QI communication graph and/or those listed in the set $\mathscr{E}$.

The QI communication graph design set

We now define a set of strongly connected and physically realizable communication graphs that generate QI subspace constraints as specified in equation (15) – in particular, the base and maximal QI graphs correspond to the boundary points of this set.

Proposition 2 *Given a plant G_{22} and a set of directed edges $\mathscr{E}$, let the adjacency matrices Γ_{base} and $\Gamma_{\max}$ of the base and maximal QI communication graphs be defined as in* (19) *and* (25)*, respectively. Then an adjacency matrix Γ corresponds to a strongly connected communication graph that can be physically built and that generates a quadratically invariant subspace $\mathcal{S}(\Gamma)$ of the form* (15) *if*

$$bsupp(\Gamma_{base}) \subseteq bsupp(\Gamma) \subseteq bsupp(\Gamma_{\max}). \tag{27}$$

Proof: Follows from Prop. 1 and definitions (25) and (26). ∎

The following corollary is then immediate.

Corollary 1 *Let Γ_1 and Γ_2 be adjacency matrices that satisfy the nesting condition* (27) *and suppose further that $bsupp(\Gamma_1) \subseteq bsupp(\Gamma_2)$. Let $\nu_\bullet$, with $\bullet \in \{base, 1, 2, \max\}$ be the closed loop norm achieved by the optimal distributed controller implemented using communication graph $\Gamma_\bullet$. Then*

$$d(\Gamma_{base}) \geq d(\Gamma_1) \geq d(\Gamma_2) \geq d(\Gamma_{\max}), \tag{28}$$

$$\mathcal{S}(\Gamma_{base}) \subseteq \mathcal{S}(\Gamma_1) \subseteq \mathcal{S}(\Gamma_2) \subseteq \mathcal{S}(\Gamma_{\max}), \tag{29}$$

and

$$\nu_{base} \geq \nu_1 \geq \nu_2 \geq \nu_{\max} \tag{30}$$

Proof: Relations (28) and (29) follow immediately from the hypotheses of the corollary and the definitions of the delays $d(\Gamma_\bullet)$ and the subspaces $\mathcal{S}(\Gamma_\bullet)$ as given in (14) and (15), respectively. The condition (30) on the norms $\nu_\bullet$ follows immediately from the subspace nesting condition (29) and the fact that the optimal norm $\nu_\bullet$ achievable by a distributed controller implemented using a communication graph with adjacency matrix $\Gamma_\bullet$ is specified by the optimal value of the objective function of the optimization problem (4) with distributed constraint $\mathcal{S}(\Gamma_\bullet)$. ■

Corollary 1 states that as more edges are added to the base QI communication graph, the performance of the optimal distributed controller implemented on the resulting communication graph improves. Thus there is a quantifiable tradeoff between the communication complexity and the closed loop performance of the resulting distributed optimal controller. To fully explore this tradeoff, the controller designer would have to enumerate the *QI communication graph design set* which is composed of adjacency matrices satisfying the nesting condition (27). Denoting this set by $\mathscr{G}$, a simple computation shows that $|\mathscr{G}| = 2^{|\mathscr{E}|}$ – thus the controller designer has to consider a set of graphs of cardinality exponential in the number of possible additional communication links. This poor scaling motivates the need for a principled approach to exploring the design space of communication graphs via the regularized optimization problem (9).

5 The Communication Graph Co-Design Algorithm

In this section we leverage Propositions 1 and 2 as well as tools from approximation theory [3], [4,5] to construct a convex penalty function $\|\cdot\|_{\text{comm}}$, which we call the *communication link norm*, that allows the controller designer to explore the QI communication graph design set $\mathscr{G}$ in a principled manner via the regularized convex optimization problem (9). We then propose a communication architecture/control law co-design algorithm based on this optimization problem and show that it indeed does produce strongly connected communication graphs that generate quadratically invariant subspaces.

5.1 The Communication Link Norm

Recall that our approach to the co-design task is to induce suitable structure in the expression $\mathfrak{C}(V^\star)$, where $V^\star$ is the solution to the regularized convex optimization problem (9) employing the yet to be specified convex penalty function $\|\cdot\|_{\text{comm}}$. We argued that the structure induced in the expression $\mathfrak{C}(V^\star)$ should correspond to a strongly connected communication graph that generates a QI subspace of the form (5), and characterized a set of graphs satisfying these properties, namely the QI communication graph design set $\mathscr{G}$. To explore the QI communication graph design set $\mathscr{G}$, we begin with the base QI communication graph Γ_{base} and augment it with additional communication links drawn from the set $\mathscr{E}$. The convex penalty function $\|\cdot\|_{\text{comm}}$ used in the regularized optimization problem (9) should therefore penalize the use of such additional communication links – in this way the controller designer can tradeoff between communication complexity and closed loop performance by varying the regularization weight λ in optimization problem (9).

We view distributed controllers implemented using a dense communication graph as being composed of a superposition of simple *atomic* controllers that are implemented using simple communication graphs, i.e., using communication graphs obtained by adding a small number of edges to

the base QI communication graph. This viewpoint suggests choosing the convex penalty function $\|\cdot\|_{\text{comm}}$ to be an atomic norm [3–5].

Indeed, if one seeks a solution $X^\star$ that can be composed as a linear combination of a small number of atoms drawn from a set $\mathscr{A}$, then a useful approach, as described in [3, 29–34], to induce such structure in the solution of an optimization problem is to employ a convex penalty function that is given by the atomic norm induced by the atoms $\mathscr{A}$ [4,5]. Examples of the types of structured solutions one may desire include sparse, group sparse and signed vectors, and low-rank, permutation and orthogonal matrices [3]. Specifically, if one desires a solution $X^\star$ that admits a decomposition of the form

$$X^\star = \sum_{i=1}^{r} c_i A_i,\ A_i \in \mathscr{A},\ c_i \geq 0 \tag{31}$$

for a set of appropriately scaled and centered atoms $\mathscr{A}$, and a small number r relative to the ambient dimension, then solving

$$\underset{X}{\text{minimize}} \quad \|\mathfrak{A}(X)\|_{\mathcal{H}_2}^2 + \lambda \|X\|_{\mathscr{A}} \tag{32}$$

with $\mathfrak{A}(\cdot)$ an affine map, and the atomic norm $\|\cdot\|_{\mathscr{A}}$ given by[8]

$$\|X\|_{\mathscr{A}} := \inf \left\{ \textstyle\sum_{A\in\mathscr{A}} c_A \,\middle|\, X = \sum_{A\in\mathscr{A}} c_A A,\ c_A \geq 0 \right\} \tag{33}$$

results in solutions that are both consistent with the data as measured in terms of the cost function $\|\mathfrak{A}(X)\|_{\mathcal{H}_2}^2$, and that admit sparse atomic decompositions, i.e., that are a combination of a small number of elements from $\mathscr{A}$.

We can therefore fully characterize our desired convex penalty function $\|\cdot\|_{\text{comm}}$ by specifying its defining atomic set $\mathscr{A}_{\text{comm}}$ and then invoking definition (33). As alluded to earlier, we choose the atoms in $\mathscr{A}_{\text{comm}}$ to correspond to distributed controllers implemented on communication graphs that can be constructed by adding a small number of communication links from the set of allowed edges $\mathscr{E}$ to the base QI communication graph Γ_{base}. In order to avoid introducing additional notation we describe the atomic set specified by communication graphs that can be constructed by adding a single communication link from the set $\mathscr{E}$ to the base QI communication graph Γ_{base} – the presented concepts then extend to the general case in a natural way. We explain why a controller designer may wish to construct an atomic set specified by more complex communication graphs in §7.

The atomic set $\mathscr{A}_{\text{comm}}$

To each communication link $(i,j) \in \mathscr{E}$ we associate the subspace $\mathcal{E}_{ij}$ given by

$$\mathcal{E}_{ij} := \mathcal{S}^{\perp}(\Gamma_{\text{base}}) \cap \mathcal{S}(\Gamma_{\text{base}} + E_{ij}). \tag{34}$$

Each subspace $\mathcal{E}_{ij}$ encodes the additional information available to the controller, relative to the base communication graph Γ_{base}, that is uniquely due to the added communication link (i,j) from sub-controller j to sub-controller i. Note that the subspaces $\mathcal{E}_{ij}$ are finite dimensional due to the strong connectedness assumption imposed on Γ_{base}, which leads to the equality $\mathcal{S}^{\perp}(\Gamma_{\text{base}}) = \mathcal{F}^{\perp}(\Gamma_{\text{base}}) \cap \mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}$.

Example 2 *Consider the base QI communication graph Γ_{base} illustrated in Figure 2 and specified by* (17)*. This communication graph generates the subspace $\mathcal{S}(\Gamma_{base})$ shown in* (18)*. We consider choosing from two additional links to augment the base communication graph Γ_{base}: a directed link*

[8] If no such decomposition exists, then $\|X\|_{\mathscr{A}} = \infty$.

from node 1 to node 3, and a directed link from node 3 to node 1. Then $\mathscr{E} = \{(1,3),(3,1)\}$ *and the corresponding subspaces* $\mathcal{E}_{ij}$ *are given by*

$$\mathcal{E}_{13} = \frac{1}{z^2}\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ * & 0 & 0 \end{bmatrix}, \mathcal{E}_{31} = \frac{1}{z^2}\begin{bmatrix} 0 & 0 & * \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

■

The atomic set is then composed of suitably normalized elements of these subspaces:

$$\mathscr{A}_{\text{comm}} := \bigcup_{(i,j)\in\mathscr{E}} \left\{ A \in \mathcal{E}_{ij} \;\middle|\; \|A\|_{\mathcal{H}_2} = 1 \right\}. \tag{35}$$

Note that we normalize our atoms relative to the $\mathcal{H}_2$ norm as this norm is isotropic; hence this normalization ensures that no atom is preferred over another within the family of atoms defined by a subspace $\mathcal{E}_{ij}$. The resulting atomic norm, which we denote the *communication link norm*, is defined on elements $X \in \mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}$ and is given by[9]

$$\begin{aligned} \|X\|_{\text{comm}} = \min_{A_{\text{base}},\{A_{ij}\}\in\mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}} \quad & \sum_{(i,j)\in\mathscr{E}} \|A_{ij}\|_{\mathcal{H}_2} \\ \text{s.t.} \quad & X = A_{\text{base}} + \sum_{(i,j)\in\mathscr{E}} A_{ij} \\ & A_{\text{base}} \in \mathcal{F}(\Gamma_{\text{base}}) \\ & A_{ij} \in \mathcal{E}_{ij} \;\forall (i,j) \in \mathscr{E}, \end{aligned} \tag{36}$$

when this optimization problem is feasible – when it is not, we set $\|X\|_{\text{comm}} = \infty$. Applying definition (36) of the communication link norm to the regularized optimization problem (9) yields the convex optimization problem

$$\begin{aligned} \underset{V,A_{\text{base}},\{A_{ij}\}\in\mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}}{\text{minimize}} \quad & \|\mathfrak{L}(V)\|_{\mathcal{H}_2}^2 + \lambda \left(\sum_{(i,j)\in\mathscr{E}} \|A_{ij}\|_{\mathcal{H}_2} \right) \\ \text{s.t.} \quad & \mathfrak{C}(V) = A_{\text{base}} + \sum_{(i,j)\in\mathscr{E}} A_{ij} \\ & A_{\text{base}} \in \mathcal{F}(\Gamma_{\text{base}}) \\ & A_{ij} \in \mathcal{E}_{ij} \;\forall (i,j) \in \mathscr{E}. \end{aligned} \tag{37}$$

Recall that in optimization problem (9) our approach to communication architecture design is to induce structure in the term $\mathfrak{C}(V)$ through the use of the communication link norm as a penalty function. Letting $\left(V^\star, \{A_{ij}^\star\}, A_{\text{base}}^\star\right)$ denote the solution to the optimization problem (37), we have that each nonzero $A_{ij}^\star$ in the atomic decomposition of $\mathfrak{C}(V)$ corresponds to an additional link from sub-controller j to sub-controller i being added to the base QI communication graph (in what follows we make precise how the structure of $\mathfrak{C}(V^\star)$ can be used to specify a communication graph). As desired, the communication link norm (36) penalizes the use of such additional links, and optimization problem (37) allows for a tradeoff between communication complexity (as measured by $\sum_{(i,j)\in\mathscr{E}} \|A_{ij}\|_{\mathcal{H}_2}$) and closed loop performance (as measured by $\|\mathfrak{L}(V)\|_{\mathcal{H}_2}^2$) of the resulting distributed controller through the regularization weight λ. Note further that $A_{\text{base}}^\star$ is not penalized by the communication link norm, ensuring that the communication graph defined by the structure of $\mathfrak{C}(V^\star)$ has Γ_{base} as a subgraph.

[9] We apply definition (33) to the components of X that lie in $\mathcal{S}^\perp(\Gamma_{\text{base}})$ to obtain an atomic norm defined on elements of that space. We then introduce an unpenalized variable $A_{\text{base}} \in \mathcal{F}(\Gamma_{\text{base}})$ to the atomic decomposition so that the resulting penalty function may be applied to elements $X \in \mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}$. The resulting penalty is actually a seminorm on $\mathcal{RH}_\infty^{\leq d(\Gamma_{\text{base}})}$ but we refer to it as a norm to maintain consistency with the terminology of [3].

Algorithm 1 Communication Architecture Co-Design

input : regularization weight λ, generalized plant G, base QI communication graph Γ_{base}, edge set $\mathscr{E}$;

output : designed communication graph adjacency matrix Γ_{des}, optimal Youla parameter $Q^\star_{\text{des}} \in \mathcal{S}(\Gamma_{\text{des}})$;

initialize:: $\Gamma_{\text{des}} \leftarrow \Gamma_{\text{base}}$, $Q^\star_{\text{des}} \leftarrow 0$;

co-design communication graph

  $\left(V^\star, \{A^\star_{ij}\}, A^\star_{\text{base}}\right) \leftarrow$ solution to optimization problem (37) with regularization weight λ;

  foreach $(i,j) \in \mathscr{E}$ *s.t.* $A^\star_{ij} \neq 0$ **do**

    $\Gamma_{\text{des}} \leftarrow \Gamma_{\text{des}} + E_{ij}$;

  end

end

refine optimal controller

  $Q^\star_{\text{des}} \leftarrow$ solution to optimization problem (7) with distributed constraint $\mathcal{F}(\Gamma_{\text{des}})$, as specified by Theorem 1;

end

return : Γ_{des}, $Q^\star_{\text{des}}$;

Remark 5 *Optimization problem* (37) *is finite dimensional, and hence can be formulated as a second order cone program by associating the finite impulse response transfer matrices* $(V, A_{base}, \{A_{ij}\})$, $\mathfrak{C}(V)$ *and* $\mathfrak{L}(V)$ *with their matrix representations. To see this, note that* $\mathcal{F}(\Gamma_{base}) \subseteq \mathcal{RH}_\infty^{\leq d(\Gamma_{base})}$, *and that by the discussion after the definition* (34) *of the subspaces* $\mathcal{E}_{ij}$, *they too satisfy* $\mathcal{E}_{ij} \subseteq \mathcal{RH}_\infty^{\leq d(\Gamma_{base})}$. *Thus the horizon* $d(\Gamma_{base})$ *over which the optimization problem* (37) *is solved is finite.*

5.2 Co-Design Algorithm and Solution Properties

In this section we formally define the communication architecture/control law co-design algorithm in terms of the optimization problem (37), and show that it can be used to co-design a strongly connected communication graph Γ that generates a QI subspace $\mathcal{S}(\Gamma)$ as defined in (15).

The co-design procedure is described in Algorithm 1. The algorithm consists of first solving the regularized optimization problem (37) to obtain solutions $\left(V^\star, \{A^\star_{ij}\}, A^\star_{\text{base}}\right)$. Using these solutions, we produce the designed communication graph Γ_{des} by augmenting the base QI communication graph Γ_{base} with all edges (i,j) such that $A^\star_{ij} \neq 0$. In particular, each non-zero term $A^\star_{ij}$ corresponds to an additional edge $(i,j) \in \mathscr{E}$ that the co-designed distributed control law will use – thus by varying the regularization weight λ the controller designer can control how much the use of an additional link is penalized by the optimization problem (37). As $\text{bsupp}\,(\Gamma_{\text{base}}) \subseteq \text{bsupp}\,(\Gamma_{\text{des}}) \subseteq \text{bsupp}\,(\Gamma_{\max})$ by construction, the designed communication graph Γ_{des} satisfies the assumptions of Proposition 2 – it is therefore strongly connected, can be physically built, and generates a subspace $\mathcal{S}(\Gamma_{\text{des}})$, according to (15), that is QI with respect to G_{22} and that admits a decomposition of the form (5). The subspace $\mathcal{S}(\Gamma_{\text{des}})$ thus satisfies the assumptions of Theorem 1, meaning that the distributed optimal controller $K^\star_{\text{des}}$ restricted to the designed subspace $\mathcal{S}(\Gamma_{\text{des}})$ is specified in terms of the solution $Q^\star_{\text{des}}$ to the convex quadratic program (7). In this way the optimal distributed controller restricted to the designed communication architecture, as well as the performance that it achieves, can be computed exactly.

Although the solution $V^\star$ to optimization problem (37) could be used to generate a distributed controller that can be implemented on the designed communication graph Γ_{des}, we claim that it

is preferable to use the solution $Q^\star_{\text{des}}$ to the non-regularized optimization problem (7). First, the use of the communication link norm penalty in the optimization problem (7) has the effect of shrinking the solution towards the origin. This means that the resulting controller specified by $V^\star$ is less aggressive, i.e., has smaller control gains, than the controller specified by the solution to the optimization problem (7) with subspace constraint $\mathcal{F}(\Gamma_{\text{des}})$.

Second, notice that for two graphs Γ_{ij} and Γ_{kl} obtained by augmenting the base QI communication graph Γ_{base} with the communication links (i,j) and (k,l), respectively, it holds that $\mathcal{S}(\Gamma_{ij}) + \mathcal{S}(\Gamma_{kl}) \subseteq \mathcal{S}(\text{bsupp}\,(\Gamma_{ij} + \Gamma_{kl}))$, with the inclusion being strict in general. In words, the linear superposition of the subspaces (15) generated by the two communications graphs Γ_{ij} and Γ_{kl} is in general a strict subset of the subspace generated by the single communication graph $\text{bsupp}\,(\Gamma_{ij} + \Gamma_{kl})$. Suppose now that the corresponding solutions $A^\star_{ij}$ and $A^\star_{kl}$ to optimization problem (37) are non-zero: then $\Gamma_{\text{des}} = \Gamma_{\text{base}} + E_{ij} + E_{kl}$, but the expression $\mathfrak{C}(V^\star)$ lies in the subspace given by $\mathcal{S}(\Gamma_{ij}) + \mathcal{S}(\Gamma_{kl})$. By the previous discussion $\mathcal{S}(\Gamma_{ij}) + \mathcal{S}(\Gamma_{kl}) \subset \mathcal{S}(\Gamma_{\text{des}})$, and thus we are imposing additional structure on the the expression $\mathfrak{C}(V^\star)$ relative to that imposed on the solution to the non-regularized optimization problem (7) with subspace constraint $\mathcal{F}(\Gamma_{\text{des}})$. This can be interpreted as the controller specified by the structure of $\mathfrak{C}(V^\star)$ not utilizing paths in the communication graph that contain both links (i,j) and (k,l). These sources of conservatism in the control law are however completely removed if one uses the solution $Q^\star_{\text{des}}$ to the non-regularized optimization problem (7).

Thus we have met our objective of developing a convex optimization based procedure for co-designing a distributed optimal controller and the communication architecture upon which it is implemented. In the next section we discuss the computational complexity of the proposed method and illustrate its efficacy on numerical examples.

6 Computational Examples

We show that the number of scalar optimization variables needed to formulate the regularized optimization problem (37) scales, up to constant factors, in a manner identical to the number of variables needed to formulate the non-regularized optimization problem (7). We then illustrate the usefulness of our approach via two examples.

Computational Complexity

We assume that the number of control inputs p_2 and the number of measurements q_2 scale as $O(n)$, where n is the number of sub-controllers in the system, i.e., we assume that there is an order constant number of actuators and sensors at each sub-controller. For an element $V \in \mathcal{RH}^{\leq d}_\infty$, each term $V^{(t)}$ in its power-series expansion is a real matrix of dimension $O(n) \times O(n)$, and thus V is defined by $O(n^2 d)$ scalar variables. The convex quadratic program (7) is therefore specified in terms of $O(n^2 d)$ variables.

To describe the number of scalar optimization variables in the regularized optimization problem (37), we need to take into account the contributions from V, A_{base} and $\{A_{ij}\}$. As per the discussion in the previous paragraph, V and A_{base} are composed of at most $O(n^2 d)$ scalar optimization variables. It can be checked that each A_{ij} has $O(d)$ optimization variables, and hence the collection $\{A_{ij}\}$ contributes $O(d|\mathscr{E}|)$ scalar optimization variables. Each sub-controller can have at most $O(n)$ additional links originating from it, and thus $|\mathscr{E}|$ scales, at worst, as $O(n^2)$. It follows that the regularized optimization problem (37) can also be specified in terms of $O(n^2 d)$ scalar optimization variables.

Finally, we note that the regularized optimization problem (37) is a second order cone program (SOCP) with at most $O(n^2d)$ second order constraints. It therefore enjoys favorable iteration complexity that scales as $O(\sqrt{d}n)$ [35], and its per-iteration complexity is at worst $O(d^3n^6)$ [36], but is typically much less when structure is exploited. In particular it is not atypical to solve a SOCP with tens to hundreds of thousands of variables [37]: noting that d scales at worst as $O(n)$, we therefore expect our method to be applicable to problems with hundreds of sub-controllers. Further, as we illustrate in the 20 sub-controller ring example below, the computational benefits of our approach compared to a brute force search are already tangible for systems with tens of sub-controllers.

6 sub-controller chain system

Consider a generalized plant (2) specified by a tridiagonal matrix $A_{\text{6-chain}} \in \mathbb{R}^{6\times 6}$ with randomly generated nonzero entries, $B_2 = C_2 = I_6$, $B_1 = C_1^\top = \begin{bmatrix} I_6 & 0_6 \end{bmatrix}$ and $D_{21} = D_{12}^\top = \begin{bmatrix} 0_6 & I_6 \end{bmatrix}$. The physical topology of the plant G_{22} is that of a 6 subsystem chain (a 3 subsystem chain is illustrated in Figure 2), and therefore the base QI communication graph $\Gamma_{\text{6-chain}} = \text{bsupp}\,(A_{\text{6-chain}})$ also defines a 6 sub-controller chain. We define the set of edges that can be added to the base graph to be

$$\mathscr{E} = \{(i,j) \in \mathscr{N} \times \mathscr{N} \,\big|\, |i-j| = 2\}, \tag{38}$$

i.e., the communication graph/control law co-design task consists of determining which additional directed communication links between second neighbors should be added to the base QI communication graph $\Gamma_{\text{6-chain}}$ to best improve the performance of the distributed optimal controller implemented on the resulting augmented communication graph.

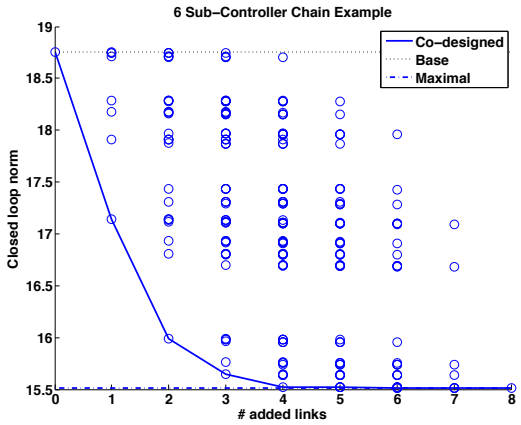

Figure 3: The closed loop norms achieved by distributed optimal controllers implemented on communication graphs constructed by adding $k = 1, \ldots, |\mathscr{E}|$ links to the base QI communication graph $\Gamma_{\text{6-chain}}$ are plotted as blue circles. The solid blue line denotes the performance achieved by distributed optimal controllers implemented on the communication graphs identified by the co-design procedure described in Algorithm 1. The dotted and dashed lines indicate the closed loop norm achieved by the distributed optimal controllers implemented on the base and maximal QI communication graphs, respectively.

In order to assess the efficacy of the proposed method in uncovering communication topologies that are well suited to distributed optimal control, we first computed the optimal closed loop performance achievable by a distributed controller implemented on every possible communication graph that can be constructed by augmenting the base QI communicating graph $\Gamma_{\text{6-chain}}$ with $k = 1, \ldots, |\mathscr{E}|$ additional links drawn from the set $\mathscr{E}$. In particular, we exhaustively explored the

QI communication graph set $\mathscr{G}$ and computed the achievable closed loop norms – these closed loop norms are plotted as blue circles in Figure 3. We then performed the co-design procedure described in Algorithm 1 for different values of regularization weight $\lambda \in [0, 50]$. The resulting closed loop norms achieved by the co-designed communication architecture/control law are plotted as a solid blue line in Figure 3. We also plot the closed loop norms achieved by controllers implemented using the base and maximal QI communication graphs.

We observe that as the regularization weight λ is increased, simpler communication topologies are generated by the co-design procedure. Further, our algorithm is able to successfully identify the optimal communication topology and the corresponding distributed optimal control law for every fixed number of additional links.

20 sub-controller ring system

Consider a generalized plant (2) specified by a matrix $A_{\text{20-ring}} \in \mathbb{R}^{20\times 20}$ with (i,j)th entry set to a nonzero randomly generated number if $|i-j| \leq 1$ where the subtraction is modulo 20 (e.g., 1-20 = 1), and 0 otherwise. The additional state-space parameters are given by $B_2 = C_2 = I_{20}$, $B_1 = C_1^\top = \begin{bmatrix} I_{20} & 0_{20} \end{bmatrix}$ and $D_{21} = D_{12}^\top = \begin{bmatrix} 0_{20} & I_{20} \end{bmatrix}$. For the example considered below, $|\lambda_{\max}(A_{\text{20-ring}})| = 2.91$. The physical topology of the plant G_{22} is that of a 20 subsystem ring, i.e., a chain topology with first and last nodes connected, and therefore the base QI communication graph $\Gamma_{\text{20-ring}} = \text{bsupp}\,(A_{\text{20-ring}})$ also defines a 20 sub-controller ring. We again define the set of edges $\mathscr{E}$ that can be added to the base graph to be those between second neighbors as in (38). In this case, the QI communication graph set $\mathscr{G}$ is too large to exhaustively explore: in particular $|\mathscr{G}| = 2^{40} \approx 10^{12}$. We performed the co-design procedure described in Algorithm 1 for different values of regularization weight $\lambda \in [0, 1000]$. The resulting closed loop norms achieved by the co-designed communication architecture/control law are plotted as a solid blue line in Figure 4. We also plot the closed loop norms achieved by controllers implemented using the base and maximal QI communication graphs. We observe again that as the regularization weight λ is increased, simpler and simpler communication topologies are designed. Notice that our method selected 10 carefully placed communication links to add to the base QI communication graph, leading to a closed loop performance only 2% higher than that achieved by the optimal controller implemented using the maximal QI communication graph.

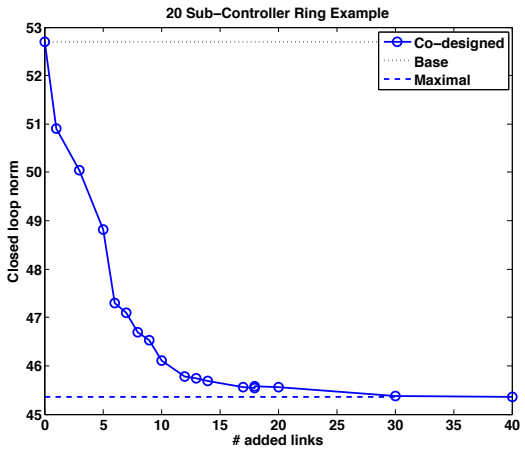

Figure 4: The solid blue line denotes the performance achieved by distributed optimal controllers implemented on the communication graphs identified by the co-design procedure described in Algorithm 1. The dotted and dashed lines indicate the closed loop norm achieved by the distributed optimal controllers implemented on the base and maximal QI communication graphs, respectively.

7 Discussion

Optimal structural recovery: It is shown in [2] that the variational solution to an $\mathcal{H}_2$ optimal control problem augmented with an atomic norm that penalizes the use of actuators can succeed in identifying an optimal actuation architecture when the dynamics of the plant satisfy certain conditions. The numerical experiments of §6 provide empirical evidence that our approach to communication architecture design identifies optimally structured controllers as well – it is of interest to see whether conditions analogous to those of [2] can provide theoretical support to the empirical success of our approach.

The k-communication link norm: The communication link norm was defined in terms of atoms corresponding to communication graphs constructed by adding a single link to the base QI communication graph. However it is possible to include atoms corresponding to communication graphs augmented with at most k-links instead, for any positive integer k; denote the resulting k-communication link norm by $\|\cdot\|_{k-\text{comm}}$. If the atoms are suitably normalized,[10] for all positive integers k_1 and k_2 satisfying $k_1 \leq k_2$ it then holds that $\|G\|_{k_1-\text{comm}} \leq \|G\|_{k_2-\text{comm}}$ for all transfer matrices G satisfying $\|G\|_{k_1-\text{comm}} < \infty$. Geometrically, restricted to the domain of $\|\cdot\|_{k_1-\text{comm}}$, the unit ball of $\|\cdot\|_{k_2-\text{comm}}$ is an inner approximation to that of $\|\cdot\|_{k_1-\text{comm}}$, and may therefore lead to simpler communication graphs when used as a penalty function in the regularized optimization problem (9). How to choose k will presumably be informed by the aforementioned conditions on optimal communication structure recovery, as well as by computational considerations, as the number of elements $\{A_{ij}\}$ required to implement the k-communciation link norm scales as $O(n^{2k})$.

Constructing base QI communication graphs: The structural assumptions made on (A, B_2, C_2) in §4 are needed to ensure that the base QI communication graph as specified in (19) is strongly connected and generates a QI subspace. However, as we note in Remark 4, any strongly connected communication topology leading to a QI subspace can be used as the base QI communication graph. Exploring how to construct base QI communication graphs in a principled way when the structural assumptions on (A, B_2, C_2) are relaxed, perhaps utilizing the methods in [38], is an interesting direction for future work. We emphasize however that the rest of the discussion in §4 remains valid once a base QI communication graph is identified even if the structural assumptions on (A, B_2, C_2) are relaxed . We also note that these issues are a consequence of the communication protocol imposed between sub-controllers – determining alternative communication protocols that allow the structural assumptions to be relaxed is also an interesting direction for future work.

Scalability: Although we expect the methods presented to be applicable to systems composed of hundreds of sub-controllers, it is important that the general approach of the RFD framework be applicable to truly large-scale systems composed of heterogeneous subsystems. The limits on the scalability of our proposed method are due to the underlying controller synthesis method [13], as opposed to being inherent to the communication link norm. To that end we have been pursuing *localized optimal control* [39] as a scalable distributed optimal controller synthesis method – an interesting direction for future work will be to see if communication architecture co-design can be incorporated into the localized optimal control framework.

Acknowledgements

The author would like to thank V. Chandrasekaran, A. Lamperski and J. C. Doyle for insightful discussions, V. D. Jonsson for her careful revisions, K. Dvijotham for references on the complexity

[10] In particular, elements $A \in \mathscr{A}_{\text{k-comm}}$ constrained to lie in a subspace $\mathcal{E}$ should be normalized as $\|A\|_{\mathcal{H}_2} = (\text{card}\,(\mathcal{E}) + \kappa)^{-\frac{1}{2}}$, where $\kappa > 0$ is a positive constant that controls how much a single atom of larger cardinality is preferred over several atoms of lower cardinality.

of SOCPs, and the anonymous reviewers for their helpful comments.

References

[1] N. Matni, "Communication delay co-design in $\mathcal{H}_2$ decentralized control using atomic norm minimization," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, Dec 2013, pp. 6522–6529.

[2] N. Matni and V. Chandrasekaran, "Regularization for design," *Submitted to the IEEE Transactions on Automatic Control*, 2015. [Online]. Available: http://arxiv.org/abs/1404.1972

[3] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational Mathematics*, vol. 12, pp. 805–849, 2012.

[4] F. Bonsall, "A general atomic decomposition theorem and Banach's closed range theorem," *The Quarterly Journal of Mathematics*, vol. 42, no. 1, pp. 9–14, 1991.

[5] G. Pisier, *Probabilistic methods in the geometry of Banach spaces*. Springer, 1986.

[6] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, 1968.

[7] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *Automatic Control, IEEE Transactions on*, vol. 51, no. 2, pp. 274–286, 2006.

[8] M. Rotkowitz, R. Cogill, and S. Lall, "Convexity of optimal control over networks with delays and arbitrary topology," *International Journal of Systems, Control and Communications*, vol. 2, no. 1-3, pp. 30–54, 2010.

[9] L. Lessard and S. Lall, "Quadratic invariance is necessary and sufficient for convexity," in *American Control Conference (ACC), 2011*, June 2011, pp. 5360–5362.

[10] B. Bamieh and P. G. Voulgaris, "A convex characterization of distributed control problems in spatially invariant systems with communication constraints," *Sys. & Control Letters*, vol. 54, no. 6, pp. 575 – 583, 2005.

[11] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, "Information structures in optimal decentralized control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, Dec 2012, pp. 1291–1306.

[12] A. Lamperski and J. Doyle, "Output feedback $\mathcal{H}_2$ model matching for decentralized systems with delays," in *American Control Conference (ACC), 2013*, June 2013, pp. 5778–5783.

[13] A. Lamperski and J. C. Doyle, "The $\mathcal{H}_2$ control problem for quadratically invariant systems with delays," *Automatic Control, IEEE Transactions on. To appear.*, 2015.

[14] P. Shah, B. N. Bhaskar, G. Tang, and B. Recht, "Linear system identification via atomic norm regularization," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 6265–6270.

[15] M. Fazel, H. Hindi, and S. Boyd, "A rank minimization heuristic with application to minimum order system approximation," in *American Control Conference (ACC), 2001*, June 2001, pp. 4734–4739.

[16] N. Matni and A. Rantzer, "Low-rank and low-order decompositions for local system identification," *CoRR*, vol. arXiv:1403.7175, 2014. [Online]. Available: http://arxiv.org/abs/1403.7175

[17] L. Ljung, "Some classical and some new ideas for identification of linear systems," *Journal of Control, Automation and Electrical Systems*, vol. 24, no. 1-2, pp. 3–10, 2013.

[18] M. Fardad, F. Lin, and M. Jovanovic, "Sparsity-promoting optimal control for a class of distributed systems," in *American Control Conference (ACC), 2011*, June 2011, pp. 2050–2055.

[19] F. Lin, M. Fardad, and M. R. Jovanovic, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *Automatic Control, IEEE Transactions on*, vol. 58, no. 9, pp. 2426–2431, 2013.

[20] V. Jonsson, A. Rantzer, and R. Murray, "A scalable formulation for engineering combination therapies for evolutionary dynamics of disease," in *American Control Conference (ACC), 2014*, June 2014, pp. 2771–2778.

[21] N. Dhingra, F. Lin, M. Fardad, and M. R. Jovanovic, "On identifying sparse representations of consensus networks," in *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems, Santa Barbara, CA*, 2012, pp. 305–310.

[22] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Distrib. Comput.*, vol. 67, no. 1, pp. 33–46, Jan. 2007.

[23] M. Fardad, F. Lin, and M. Jovanovic, "Design of optimal sparse interconnection graphs for synchronization of oscillator networks," *Automatic Control, IEEE Transactions on*, vol. 59, no. 9, pp. 2457–2462, Sept 2014.

[24] N. Matni and V. Chandrasekaran, "Regularization for design," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, Dec 2014, pp. 1111–1118.

[25] N. Dhingra, M. Jovanovic, and Z.-Q. Luo, "An ADMM algorithm for optimal sensor and actuator selection," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, Dec 2014, pp. 4039–4044.

[26] B. Polyak, M. Khlebnikov, and P. Shcherbakov, "An LMI approach to structured sparse feedback design in linear control systems," in *Control Conference (ECC), 2013 European*, July 2013, pp. 833–838.

[27] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1996.

[28] C. D. Godsil, G. Royle, and C. Godsil, *Algebraic Graph Theory*. Springer New York, 2001, vol. 207.

[29] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, Jan. 2001.

[30] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *Info. Theory, IEEE Trans. on*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[31] D. L. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ_1-norm solution is also the sparsest solution," *Comm. Pure Appl. Math*, vol. 59, pp. 797–829, 2004.

[32] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, PhD thesis, Stanford University, 2002.

[33] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, Aug. 2010.

[34] E. J. Candes and B. Recht, "Exact matrix completion via convex optimization," *Found. Comp. Math.*, vol. 9, no. 6, pp. 717–772, Dec. 2009.

[35] F. Alizadeh and D. Goldfarb, "Second-order cone programming," *Mathematical programming*, vol. 95, no. 1, pp. 3–51, 2003.

[36] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Applications of second-order cone programming," *Linear Algebra and its Applications*, vol. 284, no. 1–3, pp. 193 – 228, 1998.

[37] H. D. Mittelmann. (2014) MISOCP and large SOCP benchmark. [Online]. Available: http://plato.asu.edu/ftp/socp.html

[38] M. Rotkowitz and N. Martins, "On the nearest quadratically invariant information constraint," *Automatic Control, IEEE Transactions on*, vol. 57, no. 5, pp. 1314–1319, May 2012.

[39] Y.-S. Wang, N. Matni, and J. Doyle, "Localized LQR optimal control," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, Dec 2014, pp. 1661–1668.